We claim

1. In a network comprising a plurality of source nodes and a plurality of sink nodes, said source nodes connecting to said sink nodes by a plurality of core nodes, each of said source nodes having a multiplicity of multi-channel links to said plurality of core nodes, each of said core nodes having a multiplicity of multi-channel links to said sink nodes, a first-order smearing method for transferring data segments of a data stream from a source node to a sink nodes, each segment having a header identifying a data stream, said method applied at each source node and comprises steps of:
    (a) Logical sorting of data segments according to data stream identifiers;
    (b) Selecting a two-link path from said source node to said sink node for each data stream;
    (c) Identifying channels of a first virtual link in said two-link path;
    (d) Updating a virtual-link-specific cyclic pointer to said channels of said first virtual link; and
    (e) Assigning a data segment of said data stream to a channel indicated by said virtual-link-specific cyclic pointer.
2. The method as claimed in claim 1 wherein data received at a source node is sorted into a multiplicity of data streams.
3. The method as claimed in claim 2 wherein each of said multiplicity of data streams at a source node is defined according to a data sink.
4. The method as claimed in claim 1 wherein the data is segmented into equal-size segments.
5. The method as claimed in claim 4 wherein each of said equal-size segments includes a routing header.
6. The method as claimed in claim 5 wherein said routing header includes a sink identifier.
7. The method as claimed in claim 2 wherein the transfer of each of said data streams is rate regulated.

8. The method as claimed in claim 1 wherein the core node comprises a single-plane switching fabric having a sufficient number of input ports to support all incoming multi-channel links and a sufficient number of output ports to support all outgoing multi-channel links.
9. The method as claimed claim 8 wherein a data segment carried by any channel in any one of said incoming multi-channel links may be switched to any channel in any one of said outgoing multi-channel links.
10. In a network comprising a plurality of source nodes and a plurality of sink nodes, said source nodes connecting to said sink nodes by a plurality of core nodes, each of said source nodes having a multiplicity of multi-channel links to said plurality of core nodes, each of said core nodes having a multiplicity of multi-channel links to said sink nodes, a second-order smearing method for transferring data segments of a data stream from a source node to a sink node, each segment having a header identifying a data stream, said method applied at each source node and comprising steps of:
    (a) Logical sorting of data segments according to data stream identifiers;
    (b) Selecting a two-link path from said source node to said sink node;
    (c) Identifying channels of a virtual link;
    (d) Updating a stream-specific cyclic pointer for said data stream to said channels of said virtual link; and
    (e) Assigning a data segment of said data stream to a channel indicated by said cyclic stream-specific cyclic pointer.
11. The method as claimed in claim 10 wherein data is received at a source node is sorted into a multiplicity of data streams.
12. The method as claimed in claim 11 wherein each of said multiplicity of data streams at a source node is defined according to a data sink.
13. The method as claimed in claim 10 wherein the data is segmented into equal-size segments.
14. The method as claimed in claim 12 wherein each of said equal-size segments includes a routing header.

15. The method as claimed in claim 14 wherein said routing header includes a sink identifier.
16. The method as claimed in claim 11 wherein the transfer of each of said data streams is rate regulated.
17. The method as claimed in claim 10 wherein said second-order smearing method is implemented by said source nodes.
18. The method as claimed in claim 10 wherein the core node comprises a number of switching fabrics each having a sufficient number of input ports to support one channel from each of incoming multi-channel links and a sufficient number of output ports to support one channel in each of outgoing multi-channel links.
19. The method as claimed in claim 18 wherein each of said single-plane switching fabrics has an internal expansion.
20. A high capacity network comprising :

    (a) a plurality of source nodes each of said source nodes receiving data at a plurality of ingress ports;

    (b) a plurality of sink nodes each of said sink nodes delivering data to a plurality of egress ports;

    (c) a plurality of core nodes each of said core nodes connecting said source nodes to said sink nodes;

    (d) a data scheduler to regulate the rate of transfer of data traffic from each of said source nodes to each of said sink nodes, and

    (e) a load-balancing apparatus;

wherein each of said source nodes is connected to a subset of said core nodes by diverging multi-channel links and each of said core nodes is connected to a subset of said sink nodes by converging multi-channel links;

and wherein said data scheduler selects one of said diverging multi-channel links to carry data from a selected one of said source nodes to a selected one of said sink nodes;

and wherein said load-balancing apparatus distributes the data load equitably among the channels of each of said diverging multi-channel links and each of said converging multi-channel links.

21. The high-capacity switching node as claimed in claim 20 wherein the source/sink nodes are geographically dispersed.

22. The high-capacity switching node as claimed in claim 20 wherein the core nodes are geographically dispersed.

23. The high-capacity switching node as claimed in claim 20 wherein the multi-channel link is an optical fiber supporting wavelength-division multiplexed signals.

24. The high capacity switching node as claimed in claim 20 wherein each of said source nodes is paired with a selected one of said sink nodes and each paired source-node sink-node share a common switching fabric.

25. The high-capacity switching node as claimed in claim 20 wherein each core node comprises a plurality of parallel input-buffered space switches.

26. The high-capacity network as claimed in claim 20 wherein said load-balancing apparatus includes:

    (a) means for data sorting according to data stream identification;
    (b) means for allocating each stream to a multi-channel link leading to a core
    (c) node;
    (d) means for selecting a channel in said selected multi-channel link; and
    (e) means for assigning said each packet to said selected multi-channel link.

27. The high capacity network as claimed in claim 26 wherein the load-balancing apparatus further includes:

    (a) A segment-index memory "X";

    (b) A smearing controller; and

    (c) A segment-index memory "Y";

wherein data segments sorted according to data stream identifier are placed in said segment-index memory "X";

and wherein said smearing controller selects a selected channel in a multi-channel virtual link for a selected one of said data segments and places said selected one of said data segments in a queue in segment memory "Y" , said queue corresponding to said selected channel.

28. A high capacity switching network comprising:

(a) a plurality of source nodes each of said source nodes receiving data traffic from a plurality of ingress ports;

(b) a plurality of sink nodes each of said sink nodes delivering data traffic to a plurality of egress ports;

(c) a plurality of core nodes;

(d) a plurality of cross connectors each connecting a subset of said source nodes to a subset of said core nodes and a subset of said core nodes to a subset of said sink nodes;

(e) a data scheduler to regulate the transfer of data traffic from each of said source nodes to each of said sink nodes; and

(f) a plurality of load-balancing apparatus;

wherein each of said source nodes connects to at least one of said cross connectors by at least one multi-channel link;

and wherein each of said cross connectors connects to at least one of said core nodes by at least one multi-channel link;

and wherein each of said core nodes connects to at least one cross connector by at least one multi-channel link;

and wherein each cross connector connects to at least one source node by at least one multi-channel link;

and wherein each load-balancing apparatus is dedicated to a multi-channel virtual link to a core node through a selected one of the cross connectors to carry the data traffic received from a selected one of the source nodes to a selected one of the sink nodes to distribute the data load equitably among the channels of the selected multi-channel link.

29. The high-capacity switching node as claimed in claim 28 wherein the source/sink nodes are geographically dispersed.

30. The high-capacity switching node as claimed in claim 28 wherein the core nodes are geographically dispersed

31. The high-capacity switching node as claimed in claim 28 wherein the cross connectors are geographically dispersed.

32. The high-capacity switching node as claimed in claim 28 wherein each cross connector includes parallel planes of optical space switches, the number of planes being at least equal to the highest number of channels in a connecting link.

33. The high-capacity switching node as claimed in claim 28 wherein the connectivity of any of the cross connector is adapted to follow projected spatial traffic-intensity distributions.

34. The high-capacity network as claimed in claim 28 wherein said plurality of load-balancing apparatus includes:

(a) means for data sorting according to data stream identification;

(b) means for allocating each stream to a multi-channel link leading to a core node;

(c) means for selecting a channel in said selected multi-channel link; and

(d) means for assigning said each packet to said selected multi-channel link.

35. The high capacity network as claimed in claim 34 wherein the load-balancing apparatus further includes:

(a) a segment-index memory "X";

(b) a smearing controller; and

(c) a segment-index memory "Y";

wherein data segments sorted according to data stream identifier are placed in said segment-index memory "X";

and wherein said smearing controller selects a selected channel in a multi-channel virtual link for a selected one of said data segments and places said selected one of said data segments in a queue in segment memory "Y" , said queue corresponding to said selected channel.